



## Data Center Cabling Considerations: Point-to-Point vs Structured Cabling

The old adage that history repeats itself is very true. If we don't learn from history, we are doomed to repeat it. Many data centers today are victims of historical point-to-point cabling practices.

Direct connections - "Point-to-Point" (i.e. from switches to servers, servers to storage, servers to other servers, etc.) are problematic and costly for a variety of reasons. In the best of data center ecosystems, a standards-based structured cabling system will provide functionality and scalability with the maximum available options for current and future equipment. While Top of Rack (ToR) and End of Row (EoR) equipment mounting options are now available, these should supplement, not replace, a structured cabling system. ToR and EoR equipment placement both rely heavily on point to point cables, typically fiber jumpers and either twinax copper assemblies or stranded patch cords to connect the network or storage equipment ports to servers.

Data centers are evolving in a rather cyclical manner. When data centers (the original computer rooms) were first built, computing services were provided via a mainframe (virtualized) environment. End users' dumb terminals were connected via point to point with coax or bus cabling using twinax. Enter the PC and Intel based server platforms, and new connections were needed. We have gone through several generations of possible cabling choices: coax (thicknet, thin net), category 3, 4, 5, 5e, 6. Now, the recommended 10 Gigabit capable copper choices for a data center are category 6A, 7 and 7<sub>A</sub> channels, OM3 grade fiber for multimode capable electronics and single mode fiber for longer range electronics.

In some data centers, samples of each of these systems can still be found under the raised floor or in overhead pathways, many of which originally were point-to-point. Today however, the "from" point and "to" point are a mystery, making cable abatement (removal of abandoned cable) problematic at best. Compounding this problem was a lack of naming conventions. If the cables were labeled at both ends, the labeling may not make sense anymore. For instance, a cable may be labeled "Unix Row, Cabinet 1." Years later, the Unix row may have been replaced and new personnel may not know where the Unix row was.

There are two standards for structured cabling systems in a data center: TIA 942 and draft ISO 24764, the latter of which is slated to publish in September, 2009.


These standards were created out of need. Both data center standards have language stating that cabling should be installed to accommodate growth over the life of the data center. Moves, adds and changes for a single or a few runs are expensive compared to the same channels run as part of an overall multi-channel installation project. For the larger projects, the end user realizes benefits from project pricing, economies of scale, and lower labor rates per channel. Single channels are typically more expensive, as it is more expensive to send person-

nel to run one channel. The risk of downtime increases with continual moves, adds and changes. Pathways and spaces can be properly planned and sized up front, but can become unruly and overfilled with additional channels being added on a regular basis.



Data centers that have issues with cable plant pathways typically suffer from poor planning. Growth and new channels were added out of need without regard to

pathways. In some cases, pathways do not accommodate growth or maximum capacity over the life of the data center. Overfilled pathways cause problems with airflow, and in some cases cabling becomes deformed due to the weight load, which can adversely affect transmission properties of the channel. This is particularly true in point-to-point systems that have grown into spaghetti-like conditions over time. Likewise, data centers that have not practiced cable abatement or removal of old cabling as newer, higher performing systems are installed experience the same disheveled pathways.



**Figure 1: Top of Rack View  
- Point-to-Point Connections**

— Copper —  
— Fiber —

**Figure 1.** Depicts a ToR patching scenario between switch ports and servers without a structured cabling system. **Rack 2 to Rack 3** connections are indicative of point-to-point server-to-switch connections, also without a structured system. While proponents of these systems tout a decrease in cabling as a cost offset, further examination may negate such savings.

If a central KVM switch is used, the centralized structured cabling system would need to co-exist anyway, albeit with less channels day one. Newer electronics may have different channel minimum/maximum lengths resulting in the need for new channels. As electronics progress, the structured system may need to be added back to the data center to support future equipment choices, completely negating the savings.

It will cost more to add the structured system later as pathways, spaces, and channels were not planned for and must be installed in a live environment increasing labor costs

and the likelihood of downtime. When adding pathways and spaces, fire suppression systems and lighting may need to be moved to accommodate added overhead pathway systems. Floor voids may need to be increased and cabinets may need to be moved to allow new pathways to be routed in a non-obstructive manner for proper airflow.

Further examination highlights other disadvantages of ToR and Point-to-Point methodologies beyond the limitations outlined previously. In either the Rack 1 or Rack 2 -> Rack 3 scenario above, switch ports are dedicated to servers within a particular cabinet. This can lead to an oversubscription of ports. Suppose rack/cabinet 1 had the need for only 26 server connections for the entire rack. If a 48 port switch (ToR switching) or 48 port blade (point-to-point server to switch) is dedicated to the cabinet, this means that 22 additional ports are purchased and maintenance is being paid on those unused ports.

A greater problem occurs when the full 48 ports are used. Adding even one new server will require the purchase of another 48 port switch. In this case, assuming two network connections for the new server, an oversubscription of 46 ports will be added to the cabinet. Even in an idle state, these excess ports consume power. Two power supplies are added to the cabinet. Active maintenance and warranty costs are also associated with the additional switch and ports.

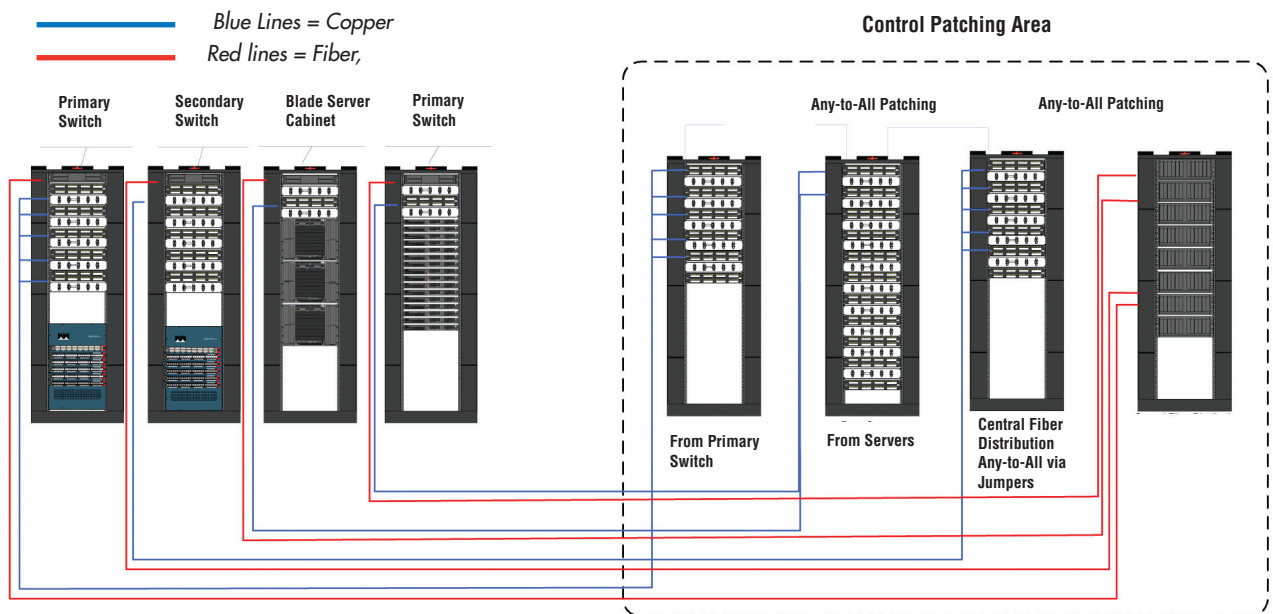
Many of these ToR technologies have limitations for cabling length. Maximum lengths range from 2-15m and are more expensive than a structured cabling channel. Short channel lengths may limit locations of equipment within the shorter cable range. With a structured cabling system, 10GBASE-T can be supported up to 100 meters of category 6A, 7 and 7A cabling and allows more options for equipment placement within the data center.

## Any-to-All Structured Cabling System

The concept behind any-to-all is quite simple. Copper and fiber panels are installed in each cabinet which correspond to copper patch panels installed in a central patching area. All fiber is run to one section of cabinets/racks in that same central patching area. This allows any equipment to be installed and connected to any other piece of equipment via either a copper patch cord or a fiber jumper. The fixed portion of the channel remains unchanged. Pathways and spaces are planned up front to properly accommodate the cabling. While this method may require more cabling up front, it has significant advantages over the life of the data center. These channels are passive and carry no reoccurring maintenance costs as realized with the addition of active electronics. If planned properly, structured cabling systems will last at least 10 years, supporting 2 or 3 generations of active electronics. The additional equipment needed for a point-to-point system will require replacement/upgrade multiple times before the structured cabling system needs to be replaced. The equipment replacement costs, not including ongoing maintenance fees, will negate any up front savings from using less cabling in a point-to-point system.

**Figure 2: Racks/Cabinets in Equipment Rows - Central Patching Area**

Example of Any-to-All Structured Cabling



The red lines (fiber connections) all arrive in the central patching area in one location. This allows any piece of equipment requiring a fiber connection to be connected to any other fiber equipment port. For instance, if a cabinet has a switch that requires a fiber connection for a SAN on day one, but needs to be changed to fiber switch connection at a later date, all that is required to connect the two ports is a fiber jumper change in the central patching area. The same is true for copper, although some data centers zone copper connections into smaller zones by function, or based on copper length and pathway requirements. As with the fiber, any copper port can be connected to any other copper port in the central patching area or within the zone.

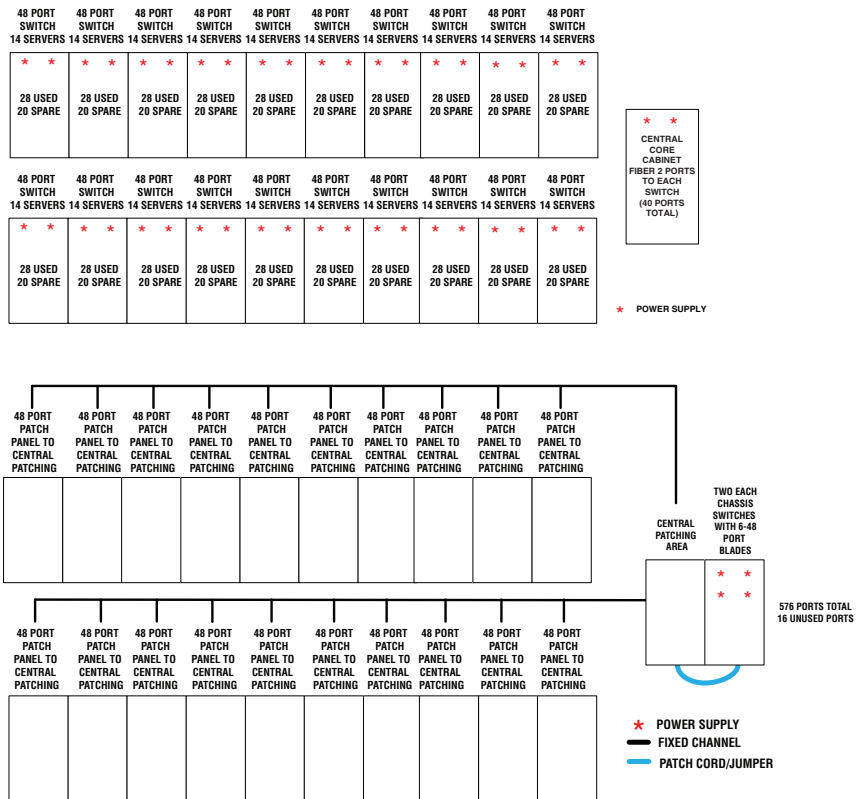
Cabling standards are written to support 2-3 generations of active electronics. An “any-to-all” configuration assures that the fixed portion of the channels is run once and remains highly unchanged if higher performing fiber and copper cabling plants are used. As a result, there will be less contractor visits to the site for MAC work as the channels already exist. Faster deployment times for equipment will be realized as no new cabling channels have to be run. They are simply connected via a patch cord. Predefined pathways and spaces will not impact cooling airflow or become overfilled as they can be properly sized for the cabling installed. Bearing in mind that the standards recommend installation of cabling accommodating growth, not only will day-one connectivity needs be supported, but also anticipated future connectivity growth needs are already accounted for.



With central patching, switch ports are not dedicated to cabinets that may not require them; therefore, active ports can be fully utilized as any port can be connected to any other port in the central patching area. Administration and documentation are enhanced as the patch panels are labeled (according to the standards) with the location at the opposite end of the channel. Patch cords and jumpers are easy to manage in cabinets rendering a more aesthetically pleasing appearance as cabinets will be tidier. In contrast, with point-to-point cabling, labeling is limited to a label attached to the end of a cable assembly.

**Figure 3: Point-to-Point Connections**

Top of the Rack view



With a structured high performing copper and fiber cabling infrastructure, recycling of cabling is minimized as several generations of electronics can utilize the same channels. Being able to utilize all switch ports lowers the number of switches and power supplies. All of these help contribute to green factors for a data center.

To further explain the power supply and switch port impact, contrasting the point-to-point, ToR scenario in section 1, in an “any-to-all” scenario, the 48 ports that would normally be dedicated to a single cabinet (ToR) can now be divided up, on demand, to any of several cabinets via the central patching area. Where autonomous LAN segments are required, VLANs or address segmentation can be used to block visibility to other segments.

**For example:** In a data center with 20 server cabinets each housing 14 servers and requiring two network connections each (560 total ports required) the port comparison is shown below.

*Note: Table assumes redundant power supplies and VLANs to segment primary and secondary networks. Counts will double if redundant switches are used.*

|                      | Number of Switches  | Number of Power Supplies (redundant) | Total Ports | Oversubscribed ports |
|----------------------|---|--------------------------------------|-------------|----------------------|
| Point-to-Point (ToR) | 20 (one 48 port switch per cabinet) 28 connections used per cab | 40                                   | 960         | 400                  |
| Central Any-to-All   | 2 chassis based with 6 ea 48 port blades                        | 4                                    | 576         | 16                   |

## Additional Power Requirements

The real limitation to equipment services within a cabinet is power. Currently in the US, the average power supplied to a cabinet is roughly 6kW<sup>1</sup> with a trend to move towards cabinets that have 18-20kW capacity. As switch ports reach full utilization, the power supplied to the cabinet may not be able to handle the load of a new server and additional switch. This may mean that new power is needed at the cabinet. A complete picture of the power required should be examined before adoption. It may not be possible from a facilities standpoint to provide enough additional power for two devices (4 power supplies in a redundant configuration). According to the Uptime Institute, one of their clients justified a \$22 million investment for new blade servers which turned into \$76 million after the necessary power and cooling capacity upgrade of \$54 million which was required for them to run.<sup>2</sup>

In "Improving Power Supply Efficiency, The Global Perspective" by Bob Mammano, Texas Instruments, "Today there are over 10 billion electronic power supplies in use worldwide, more than 3.1 billion just in the United States." Increasing the average efficiency of these power supplies by just 10% would reduce lost power by 30 billion kWhrs/year, save approximately \$3 billion per year which is equivalent to building 4 to 6 new generating plants.<sup>3</sup> Having a greater number of power supplies (as in ToR) for switches and servers will make it more difficult to upgrade to more efficient power supplies as they are introduced due to the high number of power supplies increasing replacement costs. In a collapsed scenario (central switching, central patching), fewer power supplies are needed and therefore cost less to upgrade.

Virtualization is being implemented in many data centers to decrease the number of server power supplies and to increase the operating efficiency (kW/bytes processed or IT Productivity per Embedded Watt IT-PEW) ratios within equipment. Virtualization also reduces the number of servers and the "floor space" needed to support them. This also reduces the power load to cool the room. Increasing the number of power supplies (ToR) can negate virtualization savings. Further, as servers are retired, the number of needed switch ports decreases. In a ToR configuration, this can increase the number of over-subscribed ports. In an any-to-all scenario dark fiber or non-energized copper cables may exist, but these are passive, require no power, have no reoccurring maintenance/warranty costs, and can be reused for other equipment in the future.

The efficiency of the power supply is only one power factor. To properly examine overall switch to server connections, percentage of processing load, efficiency of the power supply under various loads, required cooling, and voltage required for the overall communications must be factored into overall data center power and efficiency numbers. According to the Uptime Institute the cost to power and cool servers over the next 3 years will equal 1.5 times the price of the server hardware. Future projections extending out to 2012 show this multiplier increasing to almost 3 times even under best case assumptions, 22 times under worst case.<sup>4</sup>

Every port, network, storage, management, etc. contribute to the overall power requirements of a server. According to the US Government Data Center Energy study from Public Law 109-431 signed December 20, 2006, approximately 50% of data center power consumption is power and cooling, 29% is server consumption, and only 5% is attributed to networking equipment. The remainder is divided into storage (a highly variable factor), lighting and other systems. From a networking stand point, looking at port consumption or power draw varies greatly between various architectures (i.e. SFP+, 10GBASE-T and Fiber). Many of these reported power statistics from the manufacturers do not show the entire switch consumption, but rather make a particular architecture sound attractive by only reporting power based on consumption of an individual port, exclusive of the rest of the switch and the higher power server network interface card at the other end of the channel. For instance, a switch might report power consumption of less than 1 watt but the server NIC required can be 15-24 watts.

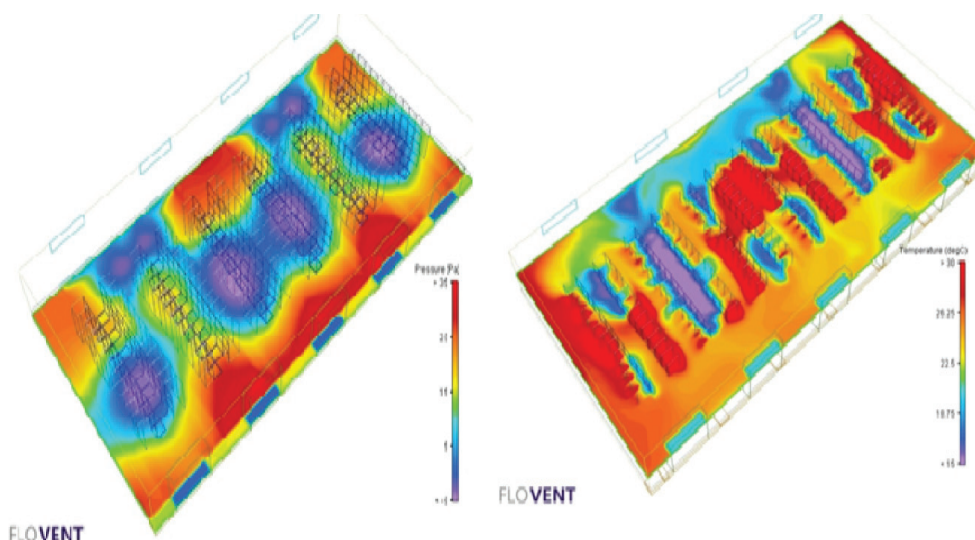
According to Kevin Tolly of the Tolly Group,<sup>5</sup> "companies that are planning for power studies and including power efficiencies in their RFP documents have difficulties in analyzing the apples to oranges comparisons in response documents. This is because numbers can be reported in a variety of ways. There has been a lack of a standard test methodology leading to our Common RFP project ([www.commonrfp.com](http://www.commonrfp.com))."<sup>5</sup> In testing at the Tolly Group, functionality in switching can vary power loads as some switches offload processing from the ASICs chips to CPU which will function at higher power. Edge switches (as those used in ToR configurations) process more instructions in CPU resulting in power spikes that may not be seen without proper testing. The goal of common RFP is to supply end users with some test methodologies to review and compare various architectures and manufacturers.

The switch port power consumption is far less, in most cases, than the server NIC at the opposite end of the channel. There has been a shift in networking led by some vendors for short point to point connections within the racks or near racks as shown in Figure 1. This shift is due in large part due to a need for 10GbE copper connections and a lack of mass manufactured low power 10GBASE-T counterparts using a structured system. The original 10GBASE-T chips had a power requirement of 10-17W per port irrespective of the switch and server power requirements. This is rapidly changing as each new version of silicon manufactured for 10GBASE-T is significantly lower power than the previous iteration. If point-to-point (currently lower power) are used for copper 10GbE communications, coexistence with a structured any-to-all system allows new technologies such as lower power 10GBASE-T to be implemented simply by installing it and connecting it via a patch cord.

End to end power and various power efficiency matrixes are provided by Tolly and The Uptime Institute amongst others. Vendor power studies may not provide a complete picture of what is required to implement the technology. Both of these groups address not only the power consumption of the device, but also the cooling required.

**Figure 3**

*Measured temperatures below the floor and at cabinet heights.*



*(illustrations provided by FloVENT)*

## Cooling Considerations

Cooling requirements are critical considerations. Poor data center equipment layout choices can cut usability by 50%.<sup>4</sup> Cooling requirements are often expressed as a function of power, but improper placement of equipment can wreak havoc on the best cooling plans. Point to point systems can land-lock equipment placement.

In Figure 3 above, we can see measured temperatures below the floor and at half cabinet heights, respectively. The ability to place equipment where it makes most sense for power and cooling can save having to purchase additional PDU whips, and in some cases, supplemental or in row cooling for hot spots. In point-to-point configurations, placement choices may be restricted to cabinets where open switch ports exist in order to avoid additional switch purchases rather than as part of the ecosystem decisions within the data center. This can lead to hot spots. Hot spots can have detrimental affects to neighboring equipment within that same cooling zone. Hot spots can be reduced with an any-to-all structured cabling system by allowing equipment to be placed where it makes the most sense for power and cooling instead of being land-locked by ToR restrictions. According to the Uptime Institute, the failure rate for equipment in the top 1/3 of the rack is 3 times greater than that of equipment at the lower 2/3's. In a structured cabling system, the passive components (cabling) are placed in the upper position leaving the cooler spaces below for the equipment. If a data center does not have enough cooling for equipment, placing the switches in a ToR position may cause them to fail prematurely due to heat as cold air supplied from under a raised floor will warm as it rises.

In conclusion, while there are several instances where point-to-point Top of Rack or End of Row connections make sense, an overall study including total equipment cost, port utilization, maintenance and power cost over time should be undertaken including both facilities and networking to make the best overall decision.

Siemon has developed several products to assist data center personnel in developing highly scalable, flexible and easy to maintain systems to support various generations of equipment singularly or in conjunction with ToR of Rack systems. Siemon's VersaPOD is an excellent example of one such innovation.

The VersaPOD™ system utilizes a central Zero-U patching zone between bayed cabinets. This space allows for any combination of copper and fiber patching and 19-inch rack-mount PDU's. Should the customer mount the switch in the top of one cabinet, the corner posts are recessed allowing cabinet to cabinet connections and allowing a switch to support multiple server cabinets increasing utilization of the switch ports. This can lower the number of switches required and save energy while providing versatile high density patching options for both copper and fiber.

For information on other Siemon innovations including category 7<sub>A</sub> TERA, Z-MAX, category 6A UTP and shielded fiber plug and play and preterminated copper and fiber trunking solutions as well as Siemon's Data Center design assistance services, please visit: [www.siemon.com](http://www.siemon.com) or contact your local Siemon representative



**Figure 4**  
VersaPOD™

## References:

<sup>1</sup> DataCenter Dynamics, Data Center Trends US, 2008

<sup>2</sup> Data Center Energy Efficiency and Productivity, Kenneth G. Brill, ([www.uptimeinstitute.com](http://www.uptimeinstitute.com))

<sup>3</sup> Power Supply Efficiency, The Global Perspective" by Bob Mammano, Texas Instruments

<sup>4</sup> The Economic Meltdown of Moore's Law, The Uptime Institute ([www.uptimeinstitute.com](http://www.uptimeinstitute.com))

<sup>5</sup> [www.tolly.com](http://www.tolly.com) and [www.commonRFP.com](http://www.commonRFP.com)

<sup>6</sup> [www.siemon.com/us/versapod](http://www.siemon.com/us/versapod) and [www.siemon.com](http://www.siemon.com)